09-27-01 A

JC760 U.S. PTO 09/26/00

Please type a plus sign (+) inside this box → [+]

PTO/SB/05 (1/98)
Approved for use through 09/30/2000. OMB 0651-0032
Patent and Trademark Office: U.S. DEPARTMENT OF COMMERCE
Under the Paperwork Reduction Act of 1995, no persons are required to respond to a collection of information unless it displays a valid OMB control number.

UTILITY PATENT APPLICATION TRANSMITTAL

(Only for new nonprovisional applications under 37 CFR 1.53(b))

Attorney Docket No.	042390.P9241
First Inventor or Application Identifier	Jacob K. Gotwals
Title	Method And System For Automatically Interpreting Computer.
Express Mail Label No.	EL485757491US

APPLICATION ELEMENTS

See MPEP chapter 600 concerning utility patent application contents.

ADDRESS TO: Assistant Commissioner for Patents, Box Patent Application, Washington, DC 20231

1. [X] * Fee Transmittal Form (e.g., PTO/SB/17) (Submit an original, and a duplicate for fee processing)
2. [X] Specification [Total Pages 25] (preferred arrangement set forth below)
   - Descriptive title of the Invention
   - Cross References to Related Applications
   - Statement Regarding Fed sponsored R & D
   - Reference to Microfiche Appendix
   - Background of the Invention
   - Brief Summary of the Invention
   - Brief Description of the Drawings (if filed)
   - Detailed Description
   - Claim(s)
   - Abstract of the Disclosure
3. [X] Drawing(s) (35 U.S.C. 113) [Total Sheets 5]
4. Oath or Declaration [Total Pages 5]
   a. [X] Newly executed (original or copy)
   b. [ ] Copy from a prior application (37 C.F.R. § 1.63(d)) (for continuation/divisional with Box 17 completed) [Note Box 5 below]
      i. [ ] DELETION OF INVENTOR(S) Signed statement attached deleting inventor(s) named in the prior application, see 37 C.F.R. §§ 1.63(d)(2) and 1.33(b).
5. [ ] Incorporation By Reference (useable if Box 4b is checked) The entire disclosure of the prior application, from which a copy of the oath or declaration is supplied under Box 4b, is considered to be part of the disclosure of the accompanying application and is hereby incorporated by reference therein.
6. [ ] Microfiche Computer Program (Appendix)
7. Nucleotide and/or Amino Acid Sequence Submission (if applicable, all necessary)
   a. [ ] Computer Readable Copy
   b. [ ] Paper Copy (identical to computer copy)
   c. [ ] Statement verifying identity of above copies

JC814 U.S. PTO 09/669604 09/26/00

ACCOMPANYING APPLICATION PARTS

8. [X] Assignment Papers (cover sheet & document(s))
9. [ ] 37 C.F.R.§3.73(b) Statement (when there is an assignee) [ ] Power of Attorney
10. [ ] English Translation Document (if applicable)
11. [X] Information Disclosure Statement (IDS)/PTO-1449 [X] Copies of IDS Citations
12. [ ] Preliminary Amendment
13. [X] Return Receipt Postcard (MPEP 503) (Should be specifically itemized)
14. [ ] * Small Entity Statement(s) (PTO/SB/09-12) [ ] Statement filed in prior application, Status still proper and desired
15. [ ] Certified Copy of Priority Document(s) (if foreign priority is claimed)
16. [ ] Other: .....

* A new statement is required to be entitled to pay small entity fees, except where one has been filed in a prior application and is being relied upon.

17. If a CONTINUING APPLICATION, check appropriate box, and supply the requisite information below and in a preliminary amendment:
    [ ] Continuation [ ] Divisional [ ] Continuation-in-part (CIP) of prior application No: _____/_____
    Prior application information: Examiner _____ Group / Art Unit: _____

18. CORRESPONDENCE ADDRESS

[ ] Customer Number or Bar Code Label (Insert Customer No. or Attach bar code label here) or [X] Correspondence address below

Name	John Travis Blakely, Sokoloff, Taylor & Zafman				
Address	12400 Wilshire Boulevard Seventh Floor				
City	Los Angeles	State	CA	Zip Code	90025-1026
Country	USA	Telephone	512-330-0844	Fax	512-330-0476

Name (Print/Type)	John Travis	Registration No. (Attorney/Agent)	43,203
Signature	John Travis	Date	9-26-00

Burden Hour Statement: This form is estimated to take 0.2 hours to complete. Time will vary depending upon the needs of the individual case. Any comments on the amount of time you are required to complete this form should be sent to the Chief Information Officer, Patent and Trademark Office, Washington, DC 20231. DO NOT SEND FEES OR COMPLETED FORMS TO THIS ADDRESS. SEND TO: Assistant Commissioner for Patents, Box Patent Application, Washington, DC 20231.

PTO/SB/17 (12/99)
Approved for use through 09/30/2000. OMB 0651-0032
Patent and Trademark Office: U.S. DEPARTMENT OF COMMERCE

FEE TRANSMITTAL for FY 2000

Patent fees are subject to annual revision.
Small Entity payments must be supported by a small entity statement, otherwise large entity fees must be paid. See Forms PTO/SB/09-12. See 37 C.F.R. §§ 1.27 and 1.28.

TOTAL AMOUNT OF PAYMENT ($) 730.00

Complete if Known	
Application Number	Not assigned
Filing Date	Herewith
First Named Inventor	Jacob K. Gotwals, et al.
Examiner Name	Not assigned
Group / Art Unit	Not assigned
Attorney Docket No.	042390.P9241

METHOD OF PAYMENT (check one)

1. ☐ The Commissioner is hereby authorized to charge indicated fees and credit any overpayments to:

Deposit Account Number: 02-2666

Deposit Account Name: Blakely, Sokoloff, Taylor & Zafman

☒ Charge Any Additional Fee Required Under 37 CFR §§ 1.16 and 1.17

2. ☒ **Payment Enclosed:** ☒ Check ☐ Money Order ☐ Other

FEE CALCULATION

1. BASIC FILING FEE

Large Entity Fee Code	Large Entity Fee ($)	Small Entity Fee Code	Small Entity Fee ($)	Fee Description	Fee Paid
101	690	201	345	Utility filing fee	690
106	310	206	155	Design filing fee	
107	480	207	240	Plant filing fee	
108	690	208	345	Reissue filing fee	
114	150	214	75	Provisional filing fee	

SUBTOTAL (1) ($) 690

2. EXTRA CLAIM FEES

	Extra Claims		Fee from below		Fee Paid
Total Claims	20 -20** =	0	X 0	=	0
Independent Claims	3 - 3** =	3	X 0	=	0
Multiple Dependent				=	

**or number previously paid, if greater; For Reissues, see below

Large Entity Fee Code	Large Entity Fee ($)	Small Entity Fee Code	Small Entity Fee ($)	Fee Description
103	18	203	9	Claims in excess of 20
102	78	202	39	Independent claims in excess of 3
104	260	204	130	Multiple dependent claim, if not paid
109	78	209	39	** Reissue independent claims over original patent
110	18	210	9	** Reissue claims in excess of 20 and over original patent

SUBTOTAL (2) ($) 0

FEE CALCULATION (continued)

3. ADDITIONAL FEES

Large Entity Fee Code	Large Entity Fee ($)	Small Entity Fee Code	Small Entity Fee ($)	Fee Description	Fee Paid
105	130	205	65	Surcharge - late filing fee or oath	
127	50	227	25	Surcharge - late provisional filing fee or cover sheet.	
139	130	139	130	Non-English specification	
147	2,520	147	2,520	For filing a request for reexamination	
112	920*	112	920*	Requesting publication of SIR prior to Examiner action	
113	1,840*	113	1,840*	Requesting publication of SIR after Examiner action	
115	110	215	55	Extension for reply within first month	
116	380	216	190	Extension for reply within second month	
117	870	217	435	Extension for reply within third month	
118	1,360	218	680	Extension for reply within fourth month	
128	1,850	228	925	Extension for reply within fifth month	
119	300	219	150	Notice of Appeal	
120	300	220	150	Filing a brief in support of an appeal	
121	260	221	130	Request for oral hearing	
138	1,510	138	1,510	Petition to institute a public use proceeding	
140	110	240	55	Petition to revive - unavoidable	
141	1,210	241	605	Petition to revive - unintentional	
142	1,210	242	605	Utility issue fee (or reissue)	
143	430	243	215	Design issue fee	
144	580	244	290	Plant issue fee	
122	130	122	130	Petitions to the Commissioner	
123	50	123	50	Petitions related to provisional applications	
126	240	126	240	Submission of Information Disclosure Stmt	
581	40	581	40	Recording each patent assignment per property (times number of properties)	40.
146	690	246	345	Filing a submission after final rejection (37 CFR § 1.129(a))	
149	690	249	345	For each additional invention to be examined (37 CFR § 1.129(b))	
Other fee (specify)					
Other fee (specify)					

* Reduced by Basic Filing Fee Paid

SUBTOTAL (3) ($) 40.00

SUBMITTED BY					
Name (Print/Type)	John Travis	Registration No. (Attorney/Agent)	43,203	Telephone	512/330-0844
Signature	John Travis			Date	9-26-00

Complete (if applicable)

WARNING:

Information on this form may become public. Credit card information should not be included on this form. Provide credit card information and authorization on PTO-2038.

UNITED STATES PATENT APPLICATION

for

METHOD AND SYSTEM FOR AUTOMATICALLY INTERPRETING COMPUTER SYSTEM PERFORMANCE MEASUREMENTS

Inventors:

Jacob K. Gotwals
Neeraj Garg

Prepared by:

Blakely, Sokoloff, Taylor & Zafman
12400 Wilshire Boulevard
Seventh Floor
Los Angeles, California
(512) 330-0844

Docket No.: 042390.P9241

EXPRESS MAIL CERTIFICATE OF MAILING

"Express Mail" mailing label number EL485757491US Date of Deposit September 26, 2000

I hereby certify that I am causing this paper or fee to be deposited with the United States Postal Service "Express Mail Post Office to Addressee" service on the date indicated above and that this paper or fee has been addressed to the Assistant Commissioner for Patents, Washington, D. C. 20231.

Shenise Ramdeen
(Typed or printed name of person mailing paper or fee)

Shenise Ramdeen
(Signature of person mailing paper or fee)

METHOD AND SYSTEM FOR AUTOMATICALLY INTERPRETING COMPUTER SYSTEM PERFORMANCE MEASUREMENTS

BACKGROUND OF THE INVENTION

1. Field of the Invention

The invention pertains generally to analyzing computer system performance. In particular, it pertains to expert systems for interpreting computer system performance measurements.

2. Description of the Related Art

Modern computer operating systems have become quite capable and equally complex, with a great deal of interdependency between the various resources that are managed by the operating system. Such resource management can include task priority, the allocation of memory, distribution of programs and data between disk/ main memory/ cache, spooling, and many others. As a result, much effort has been put into getting the maximum performance out of a system by monitoring the system and adjusting various parameters to improve the performance parameters that are considered most important in that particular system. In a related activity, application developers conduct similar optimization efforts to maximize performance in their application programs. These optimization efforts are generically known as system tuning.

Various types of analysis systems are used to implement system tuning. Some vendors have attempted to use so called 'expert systems'. Expert systems incorporate a body of knowledge that is typically possessed by experts in a particular field, and an analytical framework that extracts the relevant portions of that knowledge to troubleshoot problems. Operators can then use those systems to analyze real-world problems and provide real-world solutions to those problems within that particular field.

Some analysis systems also use 'fuzzy logic', which is necessary when all the facts needed to accurately analyze a particular problem are not known. By examining the facts that are known, and assigning a probability to other data that isn't known for certain but is likely to be true, fuzzy logic can determine an answer that is likely to be true, and provide a confidence factor to indicate the likelihood of that answer actually being correct, given the incomplete state of the underlying assumptions.

Unfortunately, many analysis systems present their results in ways that are not easy for the human user to interpret. Numerical results from various activity monitors may be presented, but it is up to the user to determine the significance of those results and decide what should be done about them. This is especially true for analysis systems used for computer system performance tuning. Also, many such systems do not allow the user to supplement the database with "knowledge plug-ins", or supplemental knowledge, to encapsulate performance tuning expertise for specific problem domains like database tuning, graphics tuning, etc., or to upgrade the existing tuning system as relevant knowledge evolves. Because of these shortcomings, only experienced system tuners are able to make full use of such

performance tuning systems, and many software developers do not fully benefit from such tools.

BRIEF DESCRIPTION OF THE DRAWINGS

Fig. 1 depicts a system of the invention.

Fig. 2 shows a measurement interface system of the invention.

Fig. 3 shows a measurement aggregator of the invention.

Fig. 4 shows a behavior interpretation system of the invention.

Fig. 5 shows a flow chart of a process of the invention.

DETAILED DESCRIPTION OF THE INVENTION

The invention adds functionality and ease-of-use to conventional computer tuning systems. The invention uses qualitative descriptions such as 'high', 'low', 'excessive', etc., to describe quantitative measurements, and uses these qualitative descriptions to generate insights/advice and present them to the user. In particular, the invention can include such things as: 1) user extensibility, 2) insights based on behavioral properties, 3) visualization of performance insights across measurement context dimensions such as time, 4) explanation of insights based on the same encoded knowledge data that is used to generate the insights, and 5) use of the abstraction of fuzzy subsystem behavioral properties to simplify knowledge

development and insight explanation, and to permit portability of the tuning system by permitting threshold values to be adjusted without affecting other areas of the database.

It is helpful to define some of the terms used in this description, and provide examples:

Behavior - a monitored measurable activity in a computer system. Example: Cache misses.

Behavioral property - a qualitative description of a behavior in a computer system. Example: A 'high' cache miss rate.

Insight - a significant combination of behavioral properties implying a result to be brought to someone's attention. Example: "An L1 data cache bottleneck can be diagnosed when CPI is high, L1 cache misses are high and L2 cache misses are low."

Relevance - The degree to which an insight is applicable to a given set of performance measurements. This is generally expressed as a numerical value, with a higher number indicating greater relevance.

Advice - a description of what should be done in response to the situation described by the insight. Example: Reduce the L1 data cache bottleneck by reducing system parameters A, B and increasing parameters X and Y.

Applicability - A measure of how likely a given piece of advice is to resolve the associated problem. This is generally expressed as a numerical value, with a higher value indicating greater applicability.

Explanation - A description of why the insight or advice is true. The explanation can repeat the criteria used to generate the insight. Example: "There appears to be an L1

data cache bottleneck because CPI is high, L1 data cache misses are high, and L2 data cache misses are low."

Visualization - a visual presentation describing the underlying information. Example: a graph of the L1 cache miss rate.

User extensibility - capability for the user to add new knowledge to the knowledge database.

Some of the features of an embodiment of the invention are described below in more detail.

User-Extensibility

The system can have functionality to permit encoded knowledge to be extended by the user or by third party knowledge developers through the use of knowledge plug-ins. This permits the expert knowledge base to grow and evolve over time, and to be customized to the particular application.

Behaviors

The system can use the abstraction of subsystem behavioral properties within the knowledge model. For example a behavior is "L2 cache misses" and a related behavioral property is "L2 cache misses are low". (L1 and L2 can refer to first- and second-level cache memories, respectively.) Using this abstraction, the encoded knowledge to monitor behaviors can be used to interpret performance counters (to determine what behavioral properties are present, e.g. are L2 cache misses high or low) separately from the encoded knowledge used to interpret behavioral properties

(to determine what insights and advice are relevant, e.g. is a L1 cache bottleneck occurring, which may depend on whether L2 cache misses are high or low.)

Performance engineers typically use the concept of subsystem behavioral properties in the way they think about and describe strategies for solving performance engineering problems, rather that describing strategies in terms of measured parameter values such as counter values. For example, a performance engineer might say "An L1 data cache bottleneck can be diagnosed when CPI is high, L1 data cache misses are high, and L2 cache misses are low". The engineer would not generally express this knowledge directly in terms of the counters used to measure each of those properties, e.g., "an L1 data cache bottleneck can be diagnosed when counter Clockticks has value 98230, counter instructions Retired has value 83948, etc."

The use of qualitative behavioral properties allows the analysis to proceed in terms that performance engineers are familiar with.

Explanations of Analysis of Behaviors

Incorporating the concept of behavioral properties into a knowledge model leads to a more natural knowledge representation that is more easily extensible by performance engineers, and enables the system to provide better automatically generated explanations of its analyses. For example, if the system has generated the insight "there appears to be an L1 data cache bottleneck", the same encoded knowledge that was used to generate this insight can be used to generate an explanation of the insight. The explanation would read "There appears to be an L1 data cache bottleneck, because CPI is high, L1 data cache misses are high, and L2 cache misses are low". Without the abstraction of behavioral properties,

automatically generated explanations would have to refer directly to the underlying performance counters. For example: "there appears to be an L1 data cache bottleneck because counter Clockticks has value 98230, counter instructions Retired has value 83948, etc", which is less useful/understandable for users.

Prioritized insights, advice, and behaviors

As previously stated, human knowledge about performance tuning often takes the form of informal diagnostic rules such as "An L1 data cache bottleneck can be diagnosed when CPI is high and L2 cache misses are low". The use of the terms *high* and *low* implies that there are not hard-and-fast thresholds for where an L1 data cache bottleneck starts occurring if CPI is above a certain threshold. Instead of hard-and-fast thresholds, the reality is that the higher CPI and L1 data cache misses, and the lower L2 cache misses, the more likely an L1 data cache bottleneck becomes. Systems that allow for the possibility that statements like "CPI is high" may be true or false or somewhere in between (e.g., when CPI is *somewhat* high) are said to make use of fuzzy logic (as opposed to standard crisp logic).

Fuzzy logic is an attempt to mimic the way performance engineers tend to think of performance engineering problems, and the fuzzy knowledge used by the invention captures this fuzziness directly. The benefit for users is that the system can use this fuzzy knowledge representation format to deliver not only insights and advice, but also a *score* for that insight and advice indicating how relevant each insight/advice is in terms of the measured data. One concept is that just as system behaviors are fuzzy (CPI may be high or low somewhere in between) insights are fuzzy too. "An L1 data cache bottleneck is occurring" may be true or false or

somewhere in between (indicating a slight bottleneck, a severe bottleneck, etc.) An alternative to fuzzy insights is *alerts*, which are crisp - an alert is either delivered to the user or it is not (an example of an alert would be "an L1 data cache bottleneck is occurring", with no associated indication or how severe the bottleneck is). Insights/advice with an associated score are more useful than alerts because they allow multiple insights/advice to be prioritized by their relevance.

Visualization or disposition of insights, advice, and behaviors

Insights, advice, and behaviors can all be fuzzy properties (that is, there is a score associated with each), which makes it possible to generate visualizations of the changes in these scores across a contextual dimension like time. For example, generating a visualization that graphs the score of the insights "L1 data cache bottleneck" against time makes it possible for the user to quickly understand the relative severity at different times, and to relate this to what other system functions are occurring at the same time. This feature (visualizing performance problem hotspots) can complement any existing hotspot-tracking visualization features, which in conventional systems only allow the user to visualize counter hotspots and counter ratio hotspots. The visualization can be in the form of a graph of the scores of insights plotted against the associated regions of code, as defined by the instruction pointer or by some other user defined metric.

Visualization is not the only function that can be performed with these results. The results can also be provided to other processes or other systems for disposition. The results can be stored or transferred, to be subsequently visualized, analyzed, or otherwise processed by those other processes or systems.

Partial Analysis Capability

In performance engineering, one rarely has a complete set of performance measurements for the systems being tuned, since there is a time cost for obtaining measurements. Expert performance engineers are able to make a partial analysis to prioritize among choices for what to measure next. Just like an expert performance engineer, the system is able to make a partial analysis based on a partial set of measurements, reporting which insights appear to be the most relevant given the partial data, and indicating what further measurements would be most useful to validate or disprove those insights.

System Description

Fig. 1 shows an overview of a system 10 of the invention. In one embodiment, the circles in Fig. 1 can represent storage for the indicated data and/or programs, while rectangles can represent processing areas. Measurement interface 100 receives measurements of various system parameters from a measurement adapter system 105, which can be a collection of one or more devices and/or software routines for collecting the desired parameters. Measurement interface 100 provides this measurement data to measurement interpreter 200, which interprets the data based on the rules provided by behavior detection knowledge 205. Measurement interpreter 200 can provide behavior data and behavior explanation data to behavior interpreter 300, which also receives insight generation knowledge 305. Behavior interpreter 300 can then provide user interface 400 with ranked insights, advice and behaviors, and explanations of those insights and advice. User interface 400 can also

receive inputs directly from measurement interpreter 200. The user interface presents the information to the user in a suitable format, which can include text and/or graphics.

Fig. 2 shows a more detailed view of the measurement interface 100 of Fig. 1. Measurement interface 100 can use measurement adapters 111, 112 and 113, which can be some of the adapters that collectively make up measurement adapter system 105. These adapters can be comprised of hardware, software, or a combination of both. Each adaptor can receive measurement data, shown as 101, 102 and 103, and provide this data to measurement aggregator 120, which collects the combined information in a form suitable for passing on to the next stage at 121 for interpretation.

Fig. 3 shows a more detailed view of measurement interpreter 200. Behavior detection knowledge (knowledge about how to detect computer subsystem behavioral properties based on measurement data) is stored at 205. A behavioral-knowledge-to-intermediate-form translator 230 can be used at knowledge database design time to generate an intermediate representation of the behavior detection knowledge that is sent to behavior detection subsystem 270 at input 121, which can interpret the intermediate form at run time to create behavior data and the associated behavior explanations. The behavior data and behavior explanations can then be presented at 271 to the following stage for interpretation. For faster performance, the intermediate representation can also be sent to a translator 240 that converts the knowledge into a C++ representation of the behavior detection knowledge, which can then be compiled by a C++ compiler 250 to generate a compiled behavior detection subsystem 260. Using inputs received at 121, behavior detection subsystem

260 can then present the behavior data and behavior explanations at 261 to the following stage for interpretation.

Thus, at run time the measurement data from the measurement aggregator 120 gets sent at 121 to either the compiled or the interpreted version of the behavior detection subsystem (260 or 270, respectively), which detects behavioral properties of the computer's subsystems, as well as explanations of the behavioral properties in terms of the measurement data.

Fig. 4 shows a more detailed version of behavior interpreter 300. The behavioral property information 261 for behavior detection subsystem 260, or behavioral property information 271 for behavior detection subsystem 270, can be sent to an insight generator 380. Insight generation knowledge (knowledge about combinations of behavioral properties that indicate significant performance conditions the user should be notified of, as well as advice associated with those conditions) is stored at 305. The insight generation subsystem uses this knowledge to interpret the behavioral properties to generate insights, advice, and explanations of insights in terms of the computer's behavioral properties. This information can then be displayed by the user interface 400, along with the explanations of behavioral properties received at 261 and 271. The information can be provided to a user in a ranked order, based on relative importance. A user, in this context, can be a person operating the system, but can also be an interface, a program, or any other thing capable of receiving the results.

Measurement adapters 111-113 can use context information (not shown in the drawing) to determine what segment of the performance measurements to collect. The entire process above can be repeated multiple times as this context information is

varied along some context dimensions, such as time, to generate insight/advice/behavior scores parameterized by that context dimension. These parameterized scores can be used as the basis for generating visualizations of insights/advice/behaviors over a given context in the user interface 400.

Procedural Description

Fig. 5 shows a flow chart 500 of some basic steps in a procedural embodiment of the invention. At step 510, significant behavioral properties of the system are identified, based on measured properties. At step 520, one or more insights associated with the identified behavioral properties are determined. At step 530, advice associated with the insights is determined. At step 540, an explanation of the insight is generated, while at step 550 an explanation of the advice is generated. Both of these explanations can then be presented to the user. However, the explanations can also be stored or transmitted to another location.

A more detailed description of a process of one embodiment of the invention follows, including those portions that involve setting up the database.

BUILDING THE KNOWLEDGE BASE

Encoding knowledge about how significant computer subsystem behavioral properties can be inferred from performance data and context data. The knowledge can be represented as follows: each behavior is represented by 1) a behavior name (textual data describing the behavior), 2) a set of property values with corresponding property names (textual data describing the possible properties of the behavior), and 3) a mathematical expression or a program, where the mathematical expression or

program computes a behavior value given the performance data and context data, where the behavior value lies within a predefined range of a set of property values, and where a proximity of the behavior value to each property value indicates the degree to which the behavior (identified by the behavior name) has that property (identified by the corresponding property value name).

Encoding knowledge about insights that capture the significance of particular combinations of computer subsystem behavioral properties that are important. The knowledge can be represented as follows: each insight is represented by 1) an insight name (textual data briefly describing the insight), 2) an insight description (textual data describing the insight and its significance at greater length), 3) a weighted list of advice describing actions that should be taken to improve the performance of the computer system when the insight is relevant (where advice is defined below), 4) a list of constraints describing conditions under which the insight should be excluded from consideration (where constraints are defined below), and 5) a list of weighted behavioral properties identifying the conditions under which the insight is relevant, where the weight of each behavioral property indicates the importance of that property relative to the other properties. An advice can include an advice description (textual data describing actions that should be taken to improve performance) and a list of constraints describing conditions under which the advice should not be given (where constraints are defined below). A constraint consists of a category (taken from a set of available categories) and an associated category value (each category has a set of possible category values).

Identifying significant behavioral properties of the computer's subsystems, based on measured performance data. The mathematical expression or program for each behavior described above can be applied to the performance data (and any associated context data) to compute a behavior value for each behavior for which the necessary measurements and context data are available.

Analyzing the behavioral properties of the computer's subsystems to compute the relevance of each insight. This can be done by combining the contribution of each behavior associated with the insight. The contribution of each behavior can consist of the proximity of the measured value of the behavior to the expected behavioral properties under which the insight is relevant, taking into account the weight of the behavior for the insight.

Computing the relevance of each advice for a given context. This can begin with identifying those insights that fall outside a predetermined acceptable range and are therefore exceptions that can be ignored. The contribution of the identified insights associated with an advice can then be combined. The contribution of an insight to an advice can, in one embodiment, be computed by taking a square of the result of multiplying the relevance of the insight by the applicability of the advice for the insight. The combination can be carried out by taking a square root of the sum of contributions from insights. Some of the terms can be further explained in this manner:

For every $Insight_i$, there is a function $Filter_i$, which takes the input Relevance, and whose output = 1. The Relevance is considered interesting, and therefore worthy of further examination, only if it falls within a predefined range of values.

Applicability is a function stored in the database, in which a given value of $Applicability(Insight_i, Advice_j)$ falls between 0 and 1.

$Relvance_i(Insight_i)$ is a computed value of relevance for $Insight_i$.

The following formulas can be used:

Let

A_k denote the relevance of Advice k

R_i denote the relevance of Insight i

V_i denote the set of possible values of relevance for Insight i

$Filter_i = V_i \rightarrow [0,1]$ ($Filter_i$ takes the measured relevance of Insight i and returns a 0 if this measured value will not be of interest to the user, and returns a 1 otherwise.)

App_{ki} denote the applicability of Advice k when Insight i is known to be relevant

Then

$$A_k = \sqrt{\Sigma_i \ (App_{ki} * Filter_i(R_i) * R_i)^2}$$

Computing the relevance of each behavior for a given context. This can begin with a relevance filter that identifies those insights with relevance values that fall outside the 'acceptable' range and are therefore exceptions that can be ignored. The contribution of the identified insights can then be combined. The contribution of an insight to a behavior can be computed by taking a square of the result of multiplying the relevance of the insight, the proximity of the behavior's measured value to its expected value, and the weight associated with the behavior for that insight. The combination can be carried out by taking a square root of the sum of contributions from insights. Some of the terms can be further explained in this manner:

$Weight_{ij}$ is a measure of the importance of $Behavior_j$ for $Insight_i$.

Proximity is a function that returns a value between 0 and 1, providing an indicator of the difference between the Expected value and the Measured value for $Behavior_j$.

The following formulas can be used:

Let

B_j denote the relevance of Behavior j

R_i denote the relevance of Insight i

E_{ij} denote the expected value for Behavior j when Insight i is relevant

m_j denote the measured value of Behavior j

W_{ij} denote the weight, or importance, of Behavior j for Insight i

P_j denote a set of possible values for Behavior j

$Proximity_j = P_j * P_j \rightarrow [0,1]$ ($Proximity_j$ is a function that takes two values from the domain of Behavior j and returns a number between 0 and 1, where 1 indicates a perfect match between the two values, and 0 indicates no match. This function might not be linear.)

C_{ij} denote the contribution of Behavior j in computing relevance of Insight i

Then

$$C_{ij} = (W_{ij} * \text{Proximity}(E_{ij}, m_j))^2$$

$$R_i = \sqrt{\Sigma_j (C_{ij})}$$

$$B_j = \sqrt{\Sigma_i (C_{ij} * R_i)}$$

Computing the relevance of each performance measurement for a given context. This can be done by combining contributions from all behaviors whose value could be computed. The contribution of a behavior to a measurement can be computed by taking a square of the result of multiplying the derivative of the behavior with respect to the measurement (or suitable approximation when the derivative is not computable) with the relevance of the behavior. The combination can be carried out by taking a square root of the sum of contributions from behaviors.

The following formulas can be used:

Let

$C = \{C_1, C_2, \ldots C_n\}$ denote a set of possible measurements

D_i denote the measured values of C_i

S_i denote the set of possible values of D_i

P_j denote the set of possible values for Behavior j

B_j denote the relevance of Behavior j (B_j is computed using m counters, $\{C_{j1}, C_{j2}, \ldots C_{jm}\}$ using function F_j:

$F_j : S_{j1} * S_{j2} * S_{j3} \ldots * S_{jm} \rightarrow P$)

F'_{jL} denote a suitable approximation of F_j with respect to C_{jL}

Then

Relevance of $C_i = \sqrt{\Sigma_j \ (B_j * F'_{jL}(D_{j1}, D_{j2} \ldots D_{jm}))}$

EXPLANATIONS PROVIDED BY SYSTEM

Generating a textual explanation of each insight for a given context. This can be done by listing all the behaviors by name, their measured values if available, and generating a textual description by finding the closest property value name and using adjectives describing the proximity of the behavior value to the property value such as "slightly", "very", etc. For example, "There appears to be an L1 data cache bottleneck, because CPI is very high, L1 data cache misses are somewhat high, and L2 cache misses are low".

Generating a textual explanation of the relevance of each advice for a given context. This can be done by listing all the insights that each item of advice is applicable to, along with a qualitative description of the relevance score of each insight. For example, "This advice has high relevance because it is applicable under the following conditions: 'L1 data cache bottleneck' (which is currently very relevant)".

Generating a textual explanation of the relevance of each behavior for a given context, and the manner in which the behavior has been computed. This can be done by listing all the associated insight names along with qualitative descriptions of their relevance scores, and listing the measurements with which the behavior was calculated and a textual representation of the mathematical expression or program that was used to generate the value for the behavior. The textual representation of a mathematical expression can be the expression itself, the textual representation of the program can be the representation made available with the program, if any, or the name of the program. For example, "This behavior has high relevance because it is associated with the following conditions: 'L1 data cache bottleneck' (which is currently very relevant)," etc., and "This behavior was computed via the following formula" (followed by a textual representation of the mathematical expression or program used to determine the behavior).

Generating a textual explanation of the relevance of each measurement for a given context. This can be done by listing the associated behaviors' names, their computed values and their relevance scores. For example, "This measurement is very relevant because it is used to compute CPI (which is currently very relevant)".

SUMMARIZING RESULTS

Generating textual summaries of insights, advice, and behaviors for a given context. This can be done by sorting them by relevance and listing their names in sorted order along with an indication of their relevance.

Generating graphical depictions of insights, advice, and behaviors for a given context. This can be done by graphing the relationship between their relevance values and context dimensions such as time, code location, etc.

Generating a list of significant sub-contexts of a given context for a given insight, advice, or behavior. This can be done by identifying sub-contexts for which the given insight, advice, or behavior has a relevance score higher than a given threshold.

In addition to the aforementioned embodiments, the invention can be implemented as instructions stored on a machine-readable medium, which can be read and executed by at least one processor to perform the functions described herein. A machine-readable medium includes any mechanism for storing or transmitting information in a form readable by a machine (e.g., a computer). For example, a machine-readable medium can include read only memory (ROM); random access memory (RAM); magnetic disk storage media; optical storage media; flash memory devices; electrical, optical, acoustical or other form of propagated signals (e.g., carrier waves, infrared signals, digital signals, etc.), and others.

The foregoing description is intended to be illustrative and not limiting. Variations will occur to those of skill in the art. Those variations are intended to be included in the invention, which is limited only by the spirit and scope of the appended claims.

We claim:

1. A method, comprising:

identifying significant ones of behavioral properties based on measured behavioral data in a computer system;

determining an insight associated with said significant ones of behavioral properties; and

determining an advice associated with the insight.

2. The method of claim 1, further comprising encoding knowledge of computer subsystem behavioral properties based on behavioral data to provide a basis for identifying said significant ones of the behavioral properties.

3. The method of claim 1, further comprising encoding knowledge of insights based on the behavioral properties to provide a basis for determining the insight.

4. The method of claim 1, further comprising generating an explanation of the insight in terms of the behavioral properties.

5. The method of claim 1, further comprising generating an explanation of the advice.

6. The method of claim 1, wherein said identifying significant ones of behavioral properties includes identifying behavior whose value is outside a predefined range.

7. The method of claim 1, further comprising providing the insight and the advice to a user.

8. A machine-readable medium having stored thereon instructions, which when executed by a processor can cause said processor to perform:

identifying significant ones of behavioral properties based on measured behavioral data in a computer system;

determining an insight associated with said significant ones of behavioral properties; and

determining an advice associated with the insight.

9. The medium of claim 8, further comprising encoding knowledge of computer subsystem behavioral properties based on behavioral data to provide a basis for identifying said significant ones of the behavioral properties.

10. The medium of claim 8, further comprising encoding knowledge of insights based on the behavioral properties to provide a basis for determining the insight.

11. The medium of claim 8, further comprising generating an explanation of the insight in terms of the behavioral properties.

12. The medium of claim 8, further comprising generating an explanation of the advice.

13. The medium of claim 8, wherein said identifying significant ones of behavioral properties includes identifying behavior whose value is outside a predefined range.

14. The medium of claim 8, further comprising providing the insight and the advice to a user.

15. A system, comprising:

a measurement interface to receive measured values of computer system performance parameters;

a measurement interpreter coupled to the measurement interface to receive measured data from the measurement interface and to provide behavioral property data;

a behavior interpreter coupled to the measurement interpreter to receive the behavioral property data and to provide insights and explanations; and

a user interface coupled to the behavior interpreter to present the insights and explanations to a user.

16. The system of claim 15, wherein the measurement interface includes a measurement adapter.

17. The system of claim 15, wherein the measurement interpreter includes a behavior knowledge interpreter.

18. The system of claim 15, wherein the measurement interpreter includes a compiler.

19. The system of claim 15, wherein the behavior interpreter includes an insight generator.

20. The system of claim 19, wherein the insight generator includes an advice generator.

ABSTRACT OF THE DISCLOSURE

An improvement in applying and interpreting encoded knowledge about computer system performance engineering and in enabling a user to more readily understand and apply that knowledge to improve computer system performance. Textual and graphical depictions of insights, advice, behaviors, and explanations can be presented, all related to improving computer system performance. Fuzzy logic can be employed to accommodate incomplete measurement data and to present results in a qualitative format, presenting relevant data as being high, low, substantial, etc., rather than simply presenting the values of the associated counters leading to those determinations. Intermediate values used to generate the results can also be processed with qualitative rather than quantitative values.

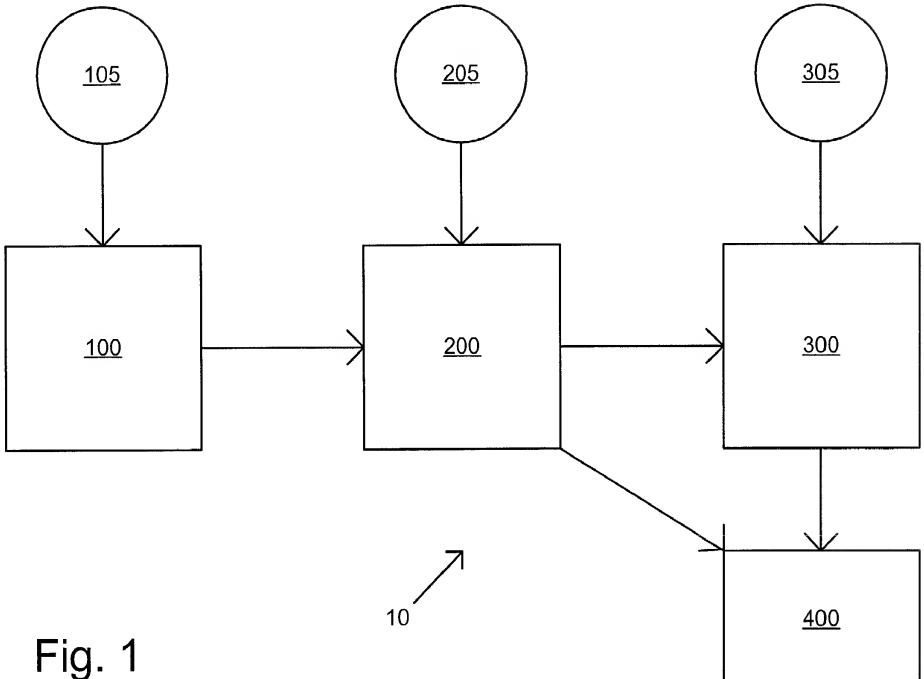

Fig. 1

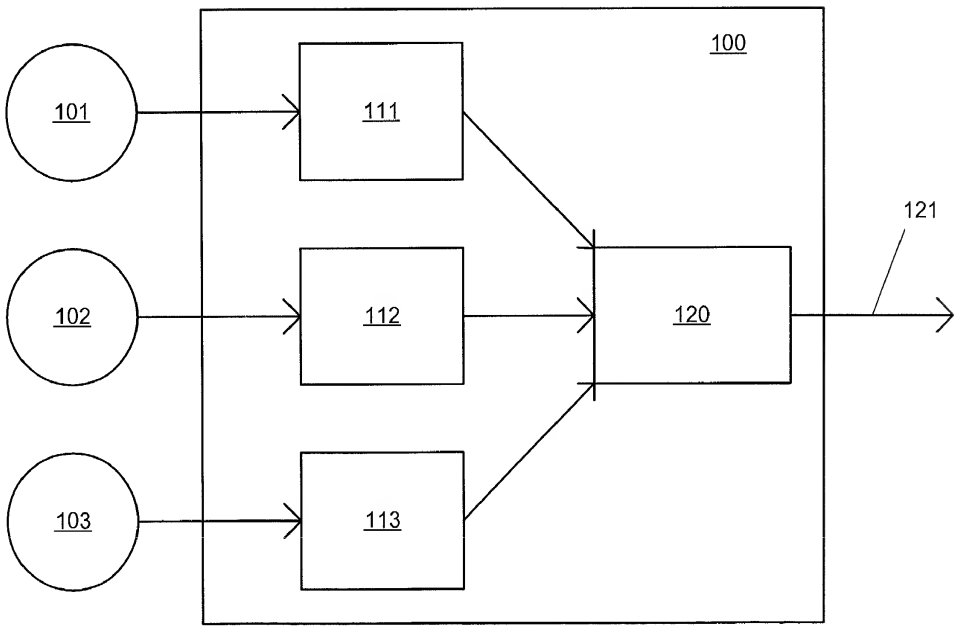

Fig. 2

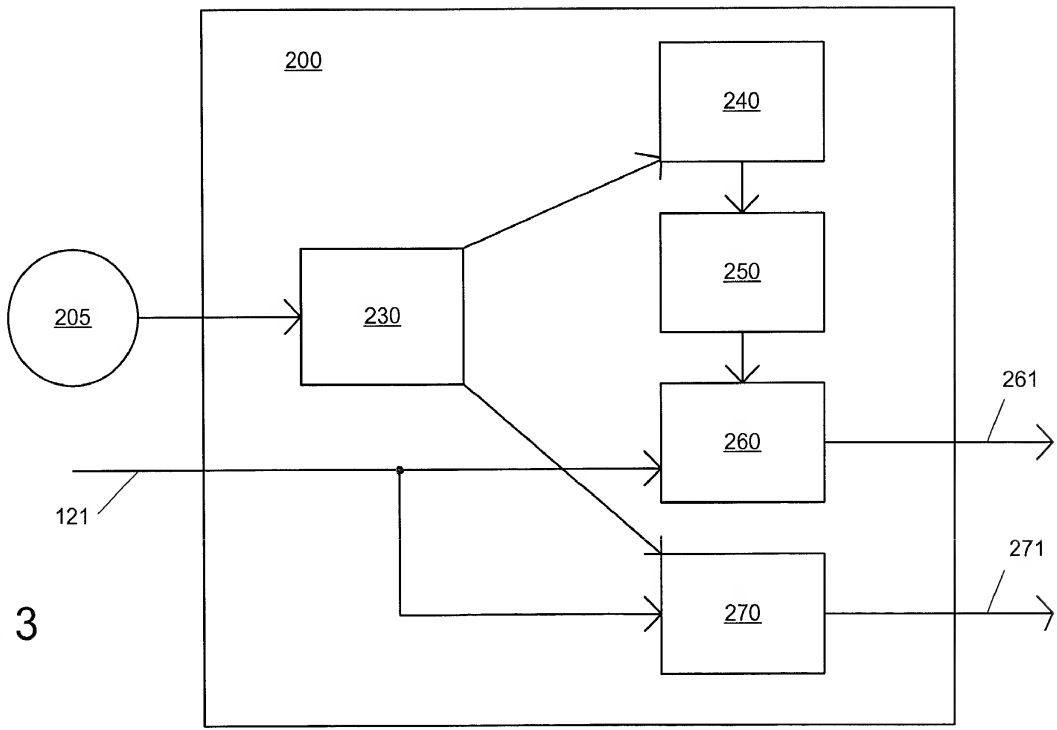

Fig. 3

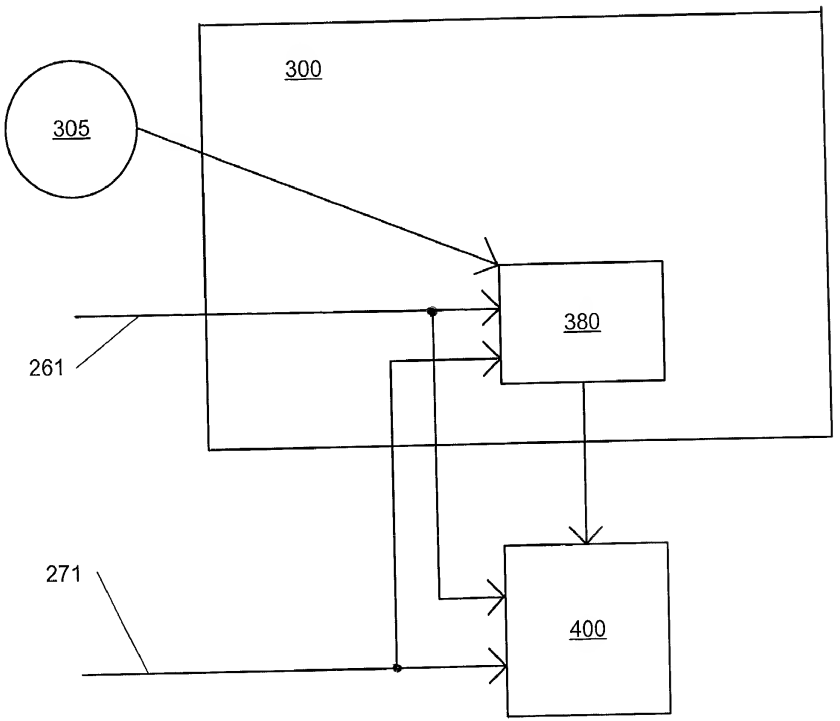

Fig. 4

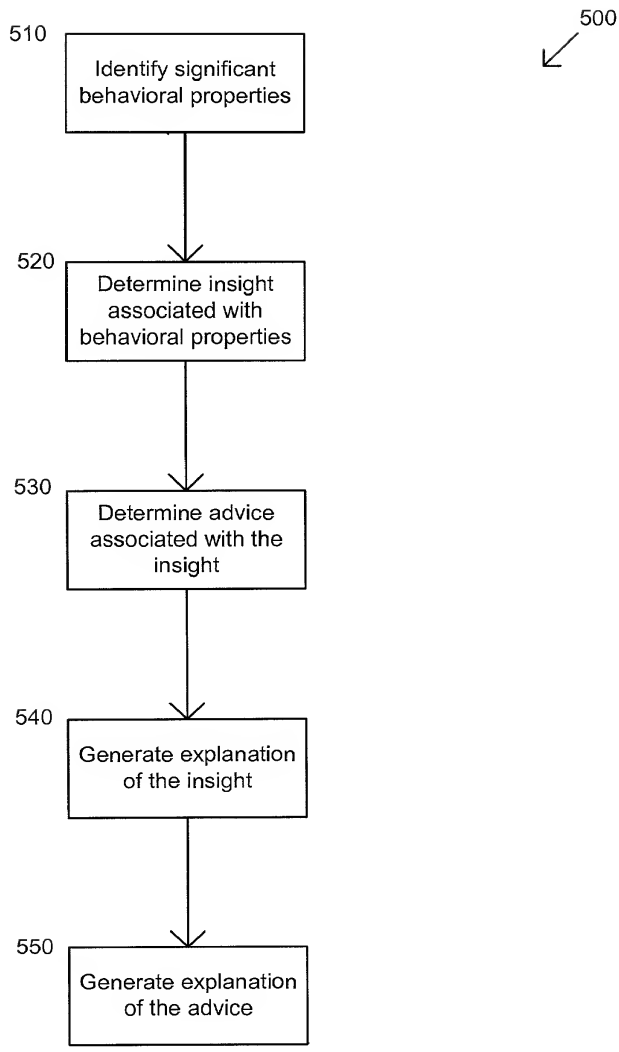
500
510
Identify significant behavioral properties
520
Determine insight associated with behavioral properties
530
Determine advice associated with the insight
540
Generate explanation of the insight
550
Generate explanation of the advice

Fig. 5

DECLARATION AND POWER OF ATTORNEY FOR PATENT APPLICATION
(FOR INTEL CORPORATION PATENT APPLICATIONS)

As a below named inventor, I hereby declare that:

My residence, post office address and citizenship are as stated below, next to my name.

I believe I am the original, first, and sole inventor (if only one name is listed below) or an original, first, and joint inventor (if plural names are listed below) of the subject matter which is claimed and for which a patent is sought on the invention entitled

METHOD AND SYSTEM FOR AUTOMATICALLY INTERPRETING COMPUTER SYSTEM PERFORMANCE MEASUREMENTS

the specification of which

_X_ is attached hereto.
___ was filed on _____ as
United States Application Number _____
or PCT International Application Number _____
and was amended on _____.
(if applicable)

I hereby state that I have reviewed and understand the contents of the above-identified specification, including the claim(s), as amended by any amendment referred to above. I do not know and do not believe that the claimed invention was ever known or used in the United States of America before my invention thereof, or patented or described in any printed publication in any country before my invention thereof or more than one year prior to this application, that the same was not in public use or on sale in the United States of America more than one year prior to this application, and that the invention has not been patented or made the subject of an inventor's certificate issued before the date of this application in any country foreign to the United States of America on an application filed by me or my legal representatives or assigns more than twelve months (for a utility patent application) or six months (for a design patent application) prior to this application.

I acknowledge the duty to disclose all information known to me to be material to patentability as defined in Title 37, Code of Federal Regulations, Section 1.56.

I hereby claim foreign priority benefits under Title 35, United States Code, Section 119(a)-(d), of any foreign application(s) for patent or inventor's certificate listed below and have also identified below any foreign application for patent or inventor's certificate having a filing date before that of the application on which priority is claimed:

Prior Foreign Application(s)			Priority Claimed	
(Number)	(Country)	(Day/Month/Year Filed)	Yes	No
(Number)	(Country)	(Day/Month/Year Filed)	Yes	No
(Number)	(Country)	(Day/Month/Year Filed)	Yes	No

I hereby claim the benefit under Title 35, United States Code, Section 119(e) of any United States provisional application(s) listed below:

Application Number	Filing Date
Application Number	Filing Date

I hereby claim the benefit under Title 35, United States Code, Section 120 of any United States application(s) listed below and, insofar as the subject matter of each of the claims of this application is not disclosed in the prior United States application in the manner provided by the first paragraph of Title 35, United States Code, Section 112, I acknowledge the duty to disclose all information known to me to be material to patentability as defined in Title 37, Code of Federal Regulations, Section 1.56 which became available between the filing date of the prior application and the national or PCT international filing date of this application:

Application Number	Filing Date	Status -- patented, pending, abandoned
Application Number	Filing Date	Status -- patented, pending, abandoned

I hereby appoint the persons listed on Appendix A hereto (which is incorporated by reference and a part of this document) as my respective patent attorneys and patent agents, with full power of substitution and revocation, to prosecute this application and to transact all business in the Patent and Trademark Office connected herewith.

Send correspondence to John Travis, Reg. No. 43,203 , BLAKELY, SOKOLOFF, TAYLOR & ZAFMAN LLP, 12400 Wilshire Boulevard 7th Floor, Los Angeles, California 90025 and direct telephone calls to John Travis, Reg. No. 43,203 , (512) 330-0844.

I hereby declare that all statements made herein of my own knowledge are true and that all statements made on information and belief are believed to be true; and further that these statements were made with the knowledge that willful false statements and the like so made are punishable by fine or imprisonment, or both, under Section 1001 of Title 18 of the United States Code and that such willful false statements may jeopardize the validity of the application or any patent issued thereon.

Full Name of Sole/First Inventor Jacob K. Gotwals

Inventor's Signature 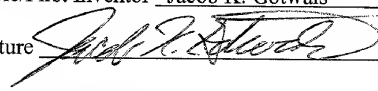 Date 9-5-2000

Residence Cupertino, California (City, State) Citizenship USA (Country)

Post Office Address 21230 Homestead Road, Apt. 122
Cupertino, California 95014

Full Name of Second/Joint Inventor Neeraj Garg

Inventor's Signature _____ Date _____

Residence San Mateo, California (City, State) Citizenship India (Country)

Post Office Address 101 Laurie Meadows Drive , Apt. 293
San Mateo, California 94403

APPENDIX A

William E. Alford, Reg. No. 37,764; Farzad E. Amini, Reg. No. P42,261; Aloysius T. C. AuYeung, Reg. No. 35,432; William Thomas Babbitt, Reg. No. 39,591; Carol F. Barry, Reg. No. 41,600; Jordan Michael Becker, Reg. No. 39,602; Bradley J. Bereznak, Reg. No. 33,474; Michael A. Bernadicou, Reg. No. 35,934; Roger W. Blakely, Jr., Reg. No. 25,831; Gregory D. Caldwell, Reg. No. 39,926; Ronald C. Card, Reg. No. P44,587; Thomas M. Coester, Reg. No. 39,637; Stephen M. De Klerk, under 37 C.F.R. § 10.9(b); Michael Anthony DeSanctis, Reg. No. 39,957; Daniel M. De Vos, Reg. No. 37,813; Robert Andrew Diehl, Reg. No. 40,992; Matthew C. Fagan, Reg. No. 37,542; Tarek N. Fahmi, Reg. No. 41,402; James Y. Go, Reg. No. 40,621; James A. Henry, Reg. No. 41,064; Willmore F. Holbrow III, Reg. No. P41,845; Sheryl Sue Holloway, Reg. No. 37,850; George W Hoover II, Reg. No. 32,992; Eric S. Hyman, Reg. No. 30,139; Dag H. Johansen, Reg. No. 36,172; William W. Kidd, Reg. No. 31,772; Erica W. Kuo, Reg. No. 42,775; Michael J. Mallie, Reg. No. 36,591; Andre L. Marais, under 37 C.F.R. § 10.9(b); Paul A. Mendonsa, Reg. No. 42,879; Darren J. Milliken, Reg. 42,004; Lisa A. Norris, Reg. No. P44,976; Chun M. Ng, Reg. No. 36,878; Thien T. Nguyen, Reg. No. 43,835; Thinh V. Nguyen, Reg. No. 42,034; Dennis A. Nicholls, Reg. No. 42,036; Daniel E. Ovanezian, Reg. No. 41,236; Babak Redjaian, Reg. No. 42,096; William F. Ryann, Reg. 44,313; James H. Salter, Reg. No. 35,668; William W. Schaal, Reg. No. 39,018; James C. Scheller, Reg. No. 31,195; Jeffrey Sam Smith, Reg. No. 39,377; Maria McCormack Sobrino, Reg. No. 31,639; Stanley W. Sokoloff, Reg. No. 25,128; Judith A. Szepesi, Reg. No. 39,393; Vincent P. Tassinari, Reg. No. 42,179; Edwin H. Taylor, Reg. No. 25,129; John F. Travis, Reg. No. 43,203; George G. C. Tseng, Reg. No. 41,355; Joseph A. Twarowski, Reg. No. 42,191; Lester J. Vincent, Reg. No. 31,460; Glenn E. Von Tersch, Reg. No. 41,364; John Patrick Ward, Reg. No. 40,216; Charles T. J. Weigell, Reg. No. 43,398; Kirk D. Williams, Reg. No. 42,229; James M. Wu, Reg. No. P45,241; Steven D. Yates, Reg. No. 42,242; and Norman Zafman, Reg. No. 26,250; my patent attorneys, and Andrew C. Chen, Reg. No. 43,544; Justin M. Dillon, Reg. No. 42,486; Paramita Ghosh, Reg. No. 42,806; and Sang Hui Kim, Reg. No. 40,450; my patent agents, of BLAKELY, SOKOLOFF, TAYLOR & ZAFMAN LLP, with offices located at 12400 Wilshire Boulevard, 7th Floor, Los Angeles, California 90025, telephone (310) 207-3800, and Alan K. Aldous, Reg. No. 31,905; Robert D. Anderson, Reg. No. 33,826; Joseph R. Bond, Reg. No. 36,458; Richard C. Calderwood, Reg. No. 35,468; Jeffrey S. Draeger, Reg. No. 41,000; Cynthia Thomas Faatz, Reg No. 39,973; Sean Fitzgerald, Reg. No. 32,027; Seth Z. Kalson, Reg. No. 40,670; David J. Kaplan, Reg. No. 41,105; Charles A. Mirho, Reg. No. 41,199; Leo V. Novakoski, Reg. No. 37,198; Naomi Obinata, Reg. No. 39,320; Thomas C. Reynolds, Reg. No. 32,488; Kenneth M. Seddon, Reg. No. 43,105; Mark Seeley, Reg. No. 32,299; Steven P. Skabrat, Reg. No. 36,279; Howard A. Skaist, Reg. No. 36,008; Steven C. Stewart, Reg. No. 33,555; Raymond J. Werner, Reg. No. 34,752; Robert G. Winkle, Reg. No. 37,474; and Charles K. Young, Reg. No. 39,435; my patent attorneys, and Thomas Raleigh Lane, Reg. No. 42,781; Calvin E. Wells; Reg. No. P43,256, Peter Lam, Reg. No. P44,855; and Gene I. Su, Reg. No. 45,140; my patent agents, of INTEL CORPORATION; and James R. Thein, Reg. No. 31,710, my patent attorney; with full power of substitution and revocation, to prosecute this application and to transact all business in the Patent and Trademark Office connected herewith.

APPENDIX B

Title 37, Code of Federal Regulations, Section 1.56
Duty to Disclose Information Material to Patentability

(a) A patent by its very nature is affected with a public interest. The public interest is best served, and the most effective patent examination occurs when, at the time an application is being examined, the Office is aware of and evaluates the teachings of all information material to patentability. Each individual associated with the filing and prosecution of a patent application has a duty of candor and good faith in dealing with the Office, which includes a duty to disclose to the Office all information known to that individual to be material to patentability as defined in this section. The duty to disclosure information exists with respect to each pending claim until the claim is cancelled or withdrawn from consideration, or the application becomes abandoned. Information material to the patentability of a claim that is cancelled or withdrawn from consideration need not be submitted if the information is not material to the patentability of any claim remaining under consideration in the application. There is no duty to submit information which is not material to the patentability of any existing claim. The duty to disclosure all information known to be material to patentability is deemed to be satisfied if all information known to be material to patentability of any claim issued in a patent was cited by the Office or submitted to the Office in the manner prescribed by §§1.97(b)-(d) and 1.98. However, no patent will be granted on an application in connection with which fraud on the Office was practiced or attempted or the duty of disclosure was violated through bad faith or intentional misconduct. The Office encourages applicants to carefully examine:

(1) Prior art cited in search reports of a foreign patent office in a counterpart application, and

(2) The closest information over which individuals associated with the filing or prosecution of a patent application believe any pending claim patentably defines, to make sure that any material information contained therein is disclosed to the Office.

(b) Under this section, information is material to patentability when it is not cumulative to information already of record or being made or record in the application, and

(1) It establishes, by itself or in combination with other information, a prima facie case of unpatentability of a claim; or

(2) It refutes, or is inconsistent with, a position the applicant takes in:

(i) Opposing an argument of unpatentability relied on by the Office, or

(ii) Asserting an argument of patentability.

A prima facie case of unpatentability is established when the information compels a conclusion that a claim is unpatentable under the preponderance of evidence, burden-of-proof standard, giving each term in the claim its broadest reasonable construction consistent with the specification, and before any consideration is given to evidence which may be submitted in an attempt to establish a contrary conclusion of patentability.

(c) Individuals associated with the filing or prosecution of a patent application within the meaning of this section are:

(1) Each inventor named in the application;

(2) Each attorney or agent who prepares or prosecutes the application; and

(3) Every other person who is substantively involved in the preparation or prosecution of the application and who is associated with the inventor, with the assignee or with anyone to whom there is an obligation to assign the application.

(d) Individuals other than the attorney, agent or inventor may comply with this section by disclosing information to the attorney, agent, or inventor.

DECLARATION AND POWER OF ATTORNEY FOR PATENT APPLICATION
(FOR INTEL CORPORATION PATENT APPLICATIONS)

As a below named inventor, I hereby declare that:

My residence, post office address and citizenship are as stated below, next to my name.

I believe I am the original, first, and sole inventor (if only one name is listed below) or an original, first, and joint inventor (if plural names are listed below) of the subject matter which is claimed and for which a patent is sought on the invention entitled

METHOD AND SYSTEM FOR AUTOMATICALLY INTERPRETING COMPUTER SYSTEM PERFORMANCE MEASUREMENTS

the specification of which

_X_ is attached hereto.
___ was filed on _____ as
United States Application Number _____
or PCT International Application Number _____
and was amended on _____.
(if applicable)

I hereby state that I have reviewed and understand the contents of the above-identified specification, including the claim(s), as amended by any amendment referred to above. I do not know and do not believe that the claimed invention was ever known or used in the United States of America before my invention thereof, or patented or described in any printed publication in any country before my invention thereof or more than one year prior to this application, that the same was not in public use or on sale in the United States of America more than one year prior to this application, and that the invention has not been patented or made the subject of an inventor's certificate issued before the date of this application in any country foreign to the United States of America on an application filed by me or my legal representatives or assigns more than twelve months (for a utility patent application) or six months (for a design patent application) prior to this application.

I acknowledge the duty to disclose all information known to me to be material to patentability as defined in Title 37, Code of Federal Regulations, Section 1.56.

I hereby claim foreign priority benefits under Title 35, United States Code, Section 119(a)-(d), of any foreign application(s) for patent or inventor's certificate listed below and have also identified below any foreign application for patent or inventor's certificate having a filing date before that of the application on which priority is claimed:

Prior Foreign Application(s)			Priority Claimed	
(Number)	(Country)	(Day/Month/Year Filed)	Yes	No
(Number)	(Country)	(Day/Month/Year Filed)	Yes	No
(Number)	(Country)	(Day/Month/Year Filed)	Yes	No

I hereby claim the benefit under Title 35, United States Code, Section 119(e) of any United States provisional application(s) listed below:

Application Number	Filing Date
Application Number	Filing Date

I hereby claim the benefit under Title 35, United States Code, Section 120 of any United States application(s) listed below and, insofar as the subject matter of each of the claims of this application is not disclosed in the prior United States application in the manner provided by the first paragraph of Title 35, United States Code, Section 112, I acknowledge the duty to disclose all information known to me to be material to patentability as defined in Title 37, Code of Federal Regulations, Section 1.56 which became available between the filing date of the prior application and the national or PCT international filing date of this application:

Application Number	Filing Date	Status -- patented, pending, abandoned
Application Number	Filing Date	Status -- patented, pending, abandoned

I hereby appoint the persons listed on Appendix A hereto (which is incorporated by reference and a part of this document) as my respective patent attorneys and patent agents, with full power of substitution and revocation, to prosecute this application and to transact all business in the Patent and Trademark Office connected herewith.

Send correspondence to John Travis, Reg. No. 43,203, BLAKELY, SOKOLOFF, TAYLOR & ZAFMAN LLP, 12400 Wilshire Boulevard 7th Floor, Los Angeles, California 90025 and direct telephone calls to John Travis, Reg. No. 43,203, (512) 330-0844.

I hereby declare that all statements made herein of my own knowledge are true and that all statements made on information and belief are believed to be true; and further that these statements were made with the knowledge that willful false statements and the like so made are punishable by fine or imprisonment, or both, under Section 1001 of Title 18 of the United States Code and that such willful false statements may jeopardize the validity of the application or any patent issued thereon.

Full Name of Sole/First Inventor Jacob K. Gotwals

Inventor's Signature _____ Date _____

Residence Cupertino, California (City, State) Citizenship USA (Country)

Post Office Address 21230 Homestead Road, Apt. 122
Cupertino, California 95014

Full Name of Second/Joint Inventor Neeraj Garg

Inventor's Signature _____ Date Sept. 15, 2000

Residence San Mateo, California (City, State) Citizenship India (Country)

Post Office Address 101 Laurie Meadows Drive , Apt. 293
San Mateo, California 94403

APPENDIX A

William E. Alford, Reg. No. 37,764; Farzad E. Amini, Reg. No. P42,261; Aloysius T. C. AuYeung, Reg. No. 35,432; William Thomas Babbitt, Reg. No. 39,591; Carol F. Barry, Reg. No. 41,600; Jordan Michael Becker, Reg. No. 39,602; Bradley J. Bereznak, Reg. No. 33,474; Michael A. Bernadicou, Reg. No. 35,934; Roger W. Blakely, Jr., Reg. No. 25,831; Gregory D. Caldwell, Reg. No. 39,926; Ronald C. Card, Reg. No. P44,587; Thomas M. Coester, Reg. No. 39,637; Stephen M. De Klerk, under 37 C.F.R. § 10.9(b); Michael Anthony DeSanctis, Reg. No. 39,957; Daniel M. De Vos, Reg. No. 37,813; Robert Andrew Diehl, Reg. No. 40,992; Matthew C. Fagan, Reg. No. 37,542; Tarek N. Fahmi, Reg. No. 41,402; James Y. Go, Reg. No. 40,621; James A. Henry, Reg. No. 41,064; Willmore F. Holbrow III, Reg. No. P41,845; Sheryl Sue Holloway, Reg. No. 37,850; George W Hoover II, Reg. No. 32,992; Eric S. Hyman, Reg. No. 30,139; Dag H. Johansen, Reg. No. 36,172; William W. Kidd, Reg. No. 31,772; Erica W. Kuo, Reg. No. 42,775; Michael J. Mallie, Reg. No. 36,591; Andre L. Marais, under 37 C.F.R. § 10.9(b); Paul A. Mendonsa, Reg. No. 42,879; Darren J. Milliken, Reg. 42,004; Lisa A. Norris, Reg. No. P44,976; Chun M. Ng, Reg. No. 36,878; Thien T. Nguyen, Reg. No. 43,835; Thinh V. Nguyen, Reg. No. 42,034; Dennis A. Nicholls, Reg. No. 42,036; Daniel E. Ovanezian, Reg. No. 41,236; Babak Redjaian, Reg. No. 42,096; William F. Ryann, Reg. 44,313; James H. Salter, Reg. No. 35,668; William W. Schaal, Reg. No. 39,018; James C. Scheller, Reg. No. 31,195; Jeffrey Sam Smith, Reg. No. 39,377; Maria McCormack Sobrino, Reg. No. 31,639; Stanley W. Sokoloff, Reg. No. 25,128; Judith A. Szepesi, Reg. No. 39,393; Vincent P. Tassinari, Reg. No. 42,179; Edwin H. Taylor, Reg. No. 25,129; John F. Travis, Reg. No. 43,203; George G. C. Tseng, Reg. No. 41,355; Joseph A. Twarowski, Reg. No. 42,191; Lester J. Vincent, Reg. No. 31,460; Glenn E. Von Tersch, Reg. No. 41,364; John Patrick Ward, Reg. No. 40,216; Charles T. J. Weigell, Reg. No. 43,398; Kirk D. Williams, Reg. No. 42,229; James M. Wu, Reg. No. P45,241; Steven D. Yates, Reg. No. 42,242; and Norman Zafman, Reg. No. 26,250; my patent attorneys, and Andrew C. Chen, Reg. No. 43,544; Justin M. Dillon, Reg. No. 42,486; Paramita Ghosh, Reg. No. 42,806; and Sang Hui Kim, Reg. No. 40,450; my patent agents, of BLAKELY, SOKOLOFF, TAYLOR & ZAFMAN LLP, with offices located at 12400 Wilshire Boulevard, 7th Floor, Los Angeles, California 90025, telephone (310) 207-3800, and Alan K. Aldous, Reg. No. 31,905; Robert D. Anderson, Reg. No. 33,826; Joseph R. Bond, Reg. No. 36,458; Richard C. Calderwood, Reg. No. 35,468; Jeffrey S. Draeger, Reg. No. 41,000; Cynthia Thomas Faatz, Reg No. 39,973; Sean Fitzgerald, Reg. No. 32,027; Seth Z. Kalson, Reg. No. 40,670; David J. Kaplan, Reg. No. 41,105; Charles A. Mirho, Reg. No. 41,199; Leo V. Novakoski, Reg. No. 37,198; Naomi Obinata, Reg. No. 39,320; Thomas C. Reynolds, Reg. No. 32,488; Kenneth M. Seddon, Reg. No. 43,105; Mark Seeley, Reg. No. 32,299; Steven P. Skabrat, Reg. No. 36,279; Howard A. Skaist, Reg. No. 36,008; Steven C. Stewart, Reg. No. 33,555; Raymond J. Werner, Reg. No. 34,752; Robert G. Winkle, Reg. No. 37,474; and Charles K. Young, Reg. No. 39,435; my patent attorneys, and Thomas Raleigh Lane, Reg. No. 42,781; Calvin E. Wells; Reg. No. P43,256, Peter Lam, Reg. No. P44,855; and Gene I. Su, Reg. No. 45,140; my patent agents, of INTEL CORPORATION; and James R. Thein, Reg. No. 31,710, my patent attorney; with full power of substitution and revocation, to prosecute this application and to transact all business in the Patent and Trademark Office connected herewith.

APPENDIX B

Title 37, Code of Federal Regulations, Section 1.56
Duty to Disclose Information Material to Patentability

(a) A patent by its very nature is affected with a public interest. The public interest is best served, and the most effective patent examination occurs when, at the time an application is being examined, the Office is aware of and evaluates the teachings of all information material to patentability. Each individual associated with the filing and prosecution of a patent application has a duty of candor and good faith in dealing with the Office, which includes a duty to disclose to the Office all information known to that individual to be material to patentability as defined in this section. The duty to disclosure information exists with respect to each pending claim until the claim is cancelled or withdrawn from consideration, or the application becomes abandoned. Information material to the patentability of a claim that is cancelled or withdrawn from consideration need not be submitted if the information is not material to the patentability of any claim remaining under consideration in the application. There is no duty to submit information which is not material to the patentability of any existing claim. The duty to disclosure all information known to be material to patentability is deemed to be satisfied if all information known to be material to patentability of any claim issued in a patent was cited by the Office or submitted to the Office in the manner prescribed by §§1.97(b)-(d) and 1.98. However, no patent will be granted on an application in connection with which fraud on the Office was practiced or attempted or the duty of disclosure was violated through bad faith or intentional misconduct. The Office encourages applicants to carefully examine:

(1) Prior art cited in search reports of a foreign patent office in a counterpart application, and

(2) The closest information over which individuals associated with the filing or prosecution of a patent application believe any pending claim patentably defines, to make sure that any material information contained therein is disclosed to the Office.

(b) Under this section, information is material to patentability when it is not cumulative to information already of record or being made or record in the application, and

(1) It establishes, by itself or in combination with other information, a prima facie case of unpatentability of a claim; or

(2) It refutes, or is inconsistent with, a position the applicant takes in:

(i) Opposing an argument of unpatentability relied on by the Office, or

(ii) Asserting an argument of patentability.

A prima facie case of unpatentability is established when the information compels a conclusion that a claim is unpatentable under the preponderance of evidence, burden-of-proof standard, giving each term in the claim its broadest reasonable construction consistent with the specification, and before any consideration is given to evidence which may be submitted in an attempt to establish a contrary conclusion of patentability.

(c) Individuals associated with the filing or prosecution of a patent application within the meaning of this section are:

(1) Each inventor named in the application;

(2) Each attorney or agent who prepares or prosecutes the application; and

(3) Every other person who is substantively involved in the preparation or prosecution of the application and who is associated with the inventor, with the assignee or with anyone to whom there is an obligation to assign the application.

(d) Individuals other than the attorney, agent or inventor may comply with this section by disclosing information to the attorney, agent, or inventor.